

CiteSeer
eprint library

Find:

Searching for PHRASE **market basket analysis visualization confidence transaction.**

Restrict to: [Header](#) [Title](#) Order by: [Citations](#) [Hubs](#) [Usage](#) [Date](#) Try: [Amazon](#) [B&N](#) [Google \(RI\)](#) [Google \(Web\)](#) [CSB](#) [DBLP](#)

No documents match Boolean query. Trying non-Boolean relevance query.

1000 documents found. Only retrieving 250 documents (System busy - maximum reduced). Retrieving documents... Order: relevance to query.

Growing and Visualizing Prediction Paths Trees in Market.. - Aria, Mola, Siciliano (Correct)

1 Growing and Visualizing Prediction Paths Trees in **Market Basket Analysis** Massimo Aria 1 Francesco and Visualizing Prediction Paths Trees in **Market Basket Analysis** Massimo Aria 1 Francesco Mola 2 Prediction Paths Trees in **Market Basket Analysis** Massimo Aria 1 Francesco Mola 2 and www.quantlet.de/scripts/compstat2002_wh/paper/full/G_01_aria.pdf

Mining Association Rules From Market Basket Data.. - Hilderman.. (1998) (Correct)

Mining Association Rules From **Market Basket** Data Using Share Measures And Mining Association Rules From **Market Basket** Data Using Share Measures And Characterized have already been selected (called itemsets)Analysis of the itemsets has enabled him to strategically www.cs.uregina.ca/~hilder/refereed_journals/ijait98.ps

Query Flocks: A Generalization of Association-Rule Mining - Tsur, Ullman.. (1998) (Correct) (37 citations)

begin our discussion by reviewing the basics of **market-basket analysis** and the a-priori algorithm for our discussion by reviewing the basics of **market-basket analysis** and the a-priori algorithm for finding theory.stanford.edu/~rajeev/postscripts/flocks.ps.gz

Dynamic Data Mining: Exploring Large Rule Spaces by Sampling - Brin, Page (1998) (Correct) (1 citation)

difficult data sets. 1 Introduction A classical **market-basket** data set consists of a large database of data sets. 1 Introduction A classical **market-basket** data set consists of a large database of cash 94]However, when standard **market basket** data analysis is applied to data sets other than market www-db.stanford.edu/~sergey/ddm.ps

Mining Association Rules with Item Constraints - Srikant, Vu, Agrawal (Correct) (77 citations)

Applications include discovering affinities for **market basket analysis** and cross-marketing, catalog include discovering affinities for **market basket analysis** and cross-marketing, catalog design, include discovering affinities for **market basket analysis** and cross-marketing, catalog design, loss-leader www.almaden.ibm.com/cs/quest/papers/kdd97_const.ps

Market Basket Analysis of Library Circulation Data - Cunningham, Frank (Correct)

Market Basket Analysis of Library Circulation Data

Market Basket Analysis of Library Circulation Data Sally Jo

but only those rules whose support and **confidence** exceed user-supplied thresholds. A rule X Y is www.cs.waikato.ac.nz/~eibe/pubs/SJ-EF-Market-Basket.ps.gz

Partial Classification using Association Rules - Ali, Manganaris, Srikant (1997) (Correct) (14 citations)

Applications include discovering affinities for **market basket analysis** and cross-marketing, catalog include discovering affinities for **market basket analysis** and cross-marketing, catalog design, include discovering affinities for **market basket analysis** and cross-marketing, catalog design, loss-leader www.almaden.ibm.com/cs/quest/papers/kdd97_class.ps

Efficient search for association rules - Webb (2000) (Correct) (6 citations)

in the most common association rule activity, **market basket analysis**, where it is often desirable to rule **analysis** is performed on domains other than **basket analysis** or when it is performed for **basket** This will often be the case when association rule **analysis** is performed on domains other than **basket** www.cm.deakin.edu.au/webb/Papers/assocrulesearch.ps

Beyond Market Baskets: Generalizing Association Rules.. - Silverstein, Brin.. (1997) (Correct) (22 citations)

Boston. Manufactured in The Netherlands. Beyond **Market Baskets**: Generalizing Association Rules to Manufactured in The Netherlands. Beyond **Market Baskets**: Generalizing Association Rules to Dependence **confidence** in the tea)coffee rule.If further **analysis** found that the dependence between coffee and tea

www-cs-students.stanford.edu/~csilvers/papers/chi2-dmkd.ps

Mining Market Basket Data Using Share Measures and .. - Hilderman.. (1998) (Correct)

Mining **Market Basket** Data Using Share Measures and
Mining **Market Basket** Data Using Share Measures and Characterized
already been selected (called itemsets 2, 14)**Analysis** of the itemsets has enabled him to strategically
www.cs.uregina.ca/~hilder/refereed_conference_proceedings/pakdd98.ps

Aspects of Network Visualization - Eick (1996) (Correct)

caused by the long transcontinental lines. 3 **Market Basket** Purchasing Correlations For many types of
by the long transcontinental lines. 3 **Market Basket** Purchasing Correlations For many types of
this perceptual tension depending on the current **analysis** task. The following examples address these
www.bell-labs.com/~eick/bibliography/1996/netvis_copyright.ps.gz

Algorithms for Association Rules - Hegland (Correct)

rule for a given data set. Having their origin in **market basked analysis**, association rules are now one of
Association rule mining originated in **market basket analysis** which aims at understanding the
data set. Having their origin in **market basked analysis**, association rules are now one of the most
datamining.anu.edu.au/.publications/2002/assoc-rules-course.pdf

Constraint-Based Rule Mining in Large, Dense Databases - Bayardo, Jr., Agrawal.. (1999) (Correct) (7 citations)

to tackle data-sets primarily from the domain of **market-basket analysis**. In **market-basket analysis**, one
data-sets primarily from the domain of **market-basket analysis**. In **market-basket analysis**, one problem
primarily from the domain of **market-basket analysis**. In **market-basket analysis**, one problem is to
www.almaden.ibm.com/cs/quest/papers/icde99_rj.ps.Z

On-Line Analytical Mining of Association Rules - Zhu (1998) (Correct) (1 citation)

among data. The discovered rules may help **market basket** or cross-sales **analysis**, decision making,
among data. The discovered rules may help **market basket** or cross-sales **analysis**, decision making, and
rules may help **market basket** or cross-sales **analysis**, decision making, and business management. In
fas.sfu.ca/pub/cs/theses/1998/HuaZhuMSc.ps.gz

Beyond Market Baskets: Generalizing Association Rules to .. - Brin, Motwani.. (1997) (Correct) (93 citations)

Beyond **Market Baskets**: Generalizing Association Rules to
Beyond **Market Baskets**: Generalizing Association Rules to Correlations
as it crosses the border. It can then do a local **analysis** of the border near the crossing. While upward
theory.stanford.edu/~rajeev/postscripts/sigmod97a.ps.gz

Dynamic Itemset Counting and Implication Rules for.. - Brin, Motwani, Ullman,.. (1997) (Correct) (152 citations)

Dynamic Itemset Counting and Implication Rules for **Market Basket** Data Sergey Brin Rajeev Motwani y
Itemset Counting and Implication Rules for **Market Basket** Data Sergey Brin Rajeev Motwani y Jeffrey
pages, and many more. We applied **market-basket analysis** to census data (see section 5)In this paper,
www-ai.cs.uni-dortmund.de/LEHRE/DATAWAREHOUSE98/Brin_etal_97a.ps.gz

Calculating a New Data Mining Algorithm for Market Basket.. - Hu, Chin, Takeichi (2000) (Correct) (1 citation)

Calculating a New Data Mining Algorithm for **Market Basket Analysis** Zhenjiang Hu 1 Wei-Ngan Chin
Calculating a New Data Mining Algorithm for **Market Basket Analysis** Zhenjiang Hu 1 Wei-Ngan Chin 2
a New Data Mining Algorithm for **Market Basket Analysis** Zhenjiang Hu 1 Wei-Ngan Chin 2 Masato
www.ipl.f.u-tokyo.ac.jp/~hu/pub/jflp01.ps.gz

Efficient Adaptive-Support Association Rule Mining for.. - Lin, Alvarez (2002) (Correct) (6 citations)

rule mining algorithms were designed with **market basket analysis** in mind. Such algorithms are
rule mining algorithms were designed with **market basket analysis** in mind. Such algorithms are inefficient
algorithms were designed with **market basket analysis** in mind. Such algorithms are inefficient for
robotics.stanford.edu/users/ronnyk/WEBKDD-DMKD/lin_alvarez_ruiz.ps

First 20 documents [Next 20](#)

Try your query at: [Amazon](#) [Barnes & Noble](#) [Google \(RI\)](#) [Google \(Web\)](#) [CSB](#) [DBLP](#)

CiteSeer - citeseer.org - [Terms of Service](#) - [Privacy Policy](#) - Copyright © 1997-2002 [NEC Research Institute](#)

CiteSeerFind: **Searching for PHRASE market basket analysis.**Restrict to: [Header](#) [Title](#) Order by: [Citations](#) [Hubs](#) [Usage](#) [Date](#) Try: [Amazon](#) [B&N](#) [Google \(RI\)](#) [Google \(Web\)](#) [CSB](#) [DBLP](#)

19 documents found. Order: citations weighted by year.

[Dynamic Itemset Counting and Implication Rules for.. - Brin, Motwani, Ullman, .. \(1997\) \(Correct\) \(152 citations\)](#)
of web pages, and many more. We applied **market-basket analysis** to census data (see section 5) In this
www-ai.cs.uni-dortmund.de/LEHRE/DATAWAREHOUSE98/Brin_etal_97a.ps.gz

One or more of the query terms is very common - only partial results have been returned. Try Google (RI).

[Query Flocks: A Generalization of Association-Rule Mining - Dick Tsur \(1998\) \(Correct\) \(34 citations\)](#)
our discussion by reviewing the basics of **market-basket analysis** and the a-priori algorithm for finding
www.cs.ucla.edu/~czdemo/tsur/flocks.ps

[Computing Iceberg Queries Efficiently - Min Fang \(1998\) \(Correct\) \(21 citations\)](#)
data warehousing, information-retrieval, **market basket analysis** in data mining, clustering and copy
www-db.stanford.edu/~shiva/Pubs/iceberg-full.ps

[Finding frequent substructures in chemical compounds - Dehaspe, Toivonen, King \(1998\) \(Correct\) \(19 citations\)](#)
A prototypical application example is in **market basket analysis**: find out which products tend to be sold
www.cs.kuleuven.ac.be/~ldh/publicaties/dehaspetoivonenking98.ps.gz

[Synopsis Data Structures for Massive Data Sets - Phillip Gibbons \(1999\) \(Correct\) \(10 citations\)](#)
fraud detection, best sellers lists, **market basket analysis**, selectivity estimation in query
www.math.tau.ac.il/~matias/papers/synopsis-soda99.ps

[The Quest Data Mining System - Agrawal, Mehta, Shafer, Srikant.. \(1996\) \(Correct\) \(27 citations\)](#)
include discovering affinities for **market basket analysis** and cross-marketing, catalog design,
www.almaden.ibm.com/u/ragrawal/papers/kdd96_quest.ps

[Mining Association Rules with Item Constraints - Srikant, Vu, Agrawal \(Correct\) \(77 citations\)](#)
include discovering affinities for **market basket analysis** and cross-marketing, catalog design,
www.almaden.ibm.com/cs/quest/papers/kdd97_const.ps

[Fast Mining of Sequential Patterns in Very Large Databases - Zaki \(1997\) \(Correct\) \(17 citations\)](#)
application is in the retail sales or **market-basket analysis** domain. Given a collection of items, a
[ftp.cs.rochester.edu/pub/papers/systems/97.tr668.Fast_mining_of_sequential_patterns_in_very_large_databases.ps](ftp://ftp.cs.rochester.edu/pub/papers/systems/97.tr668.Fast_mining_of_sequential_patterns_in_very_large_databases.ps)

[Web Mining: Pattern Discovery from World Wide Web.. - Mobasher, Jain, Han.. \(1996\) \(Correct\) \(22 citations\)](#)
to a suitable form. In particular, unlike **market basket analysis**, where a single transaction is defined
maya.cs.depaul.edu/~mobasher/papers/webminer-tr96.ps

[Partial Classification using Association Rules - Ali, Manganaris, Srikant \(1997\) \(Correct\) \(14 citations\)](#)
include discovering affinities for **market basket analysis** and cross-marketing, catalog design,
www.almaden.ibm.com/cs/quest/papers/kdd97_class.ps

[Hypergraph Based Clustering in High-Dimensional Data.. - Han, Karypis, Kumar.. \(1998\) \(Correct\) \(8 citations\)](#)
data mining applications. For example, in **market basket analysis**, a typical store sells thousands of
[ftp.cs.umn.edu/dept/users/kumar/de98-bull-cluster.ps](ftp://ftp.cs.umn.edu/dept/users/kumar/de98-bull-cluster.ps)

[Discovering Predictive Association Rules - Megiddo, Srikant \(1998\) \(Correct\) \(8 citations\)](#)
include discovering affinities for **market basket analysis** and cross-marketing, catalog design,
www.almaden.ibm.com/cs/quest/papers/kdd98_stat.ps

Stock Movement Prediction And N-Dimensional Inter-Transaction.. - Lu, al. (1998) (Correct) (3 citations)
cited application of association rules is **market basket analysis** using transaction databases from
www.cs.ust.hk/~luhj/ps/dmkd98.ps

Clustering Based On Association Rule Hypergraphs - Han, Karypis, al. (Correct) (15 citations)
common in many data mining domains (e.g. **market basket analysis**) in which the number of different items
ftp.cs.umn.edu/dept/users/kumar/cluster-hyper-pos.ps

On-Line Analytical Mining of Association Rules - Zhu (1998) (Correct) (1 citation)
.23 2.2.1 **Market basket analysis** .24
fas.sfu.ca/pub/cs/theses/1998/HuaZhuMSc.ps.gz

Towards a Cost-Effective Parallel Data Mining Approach - Zoltan Jarai (Correct)
dealt with one particular type of problem: **market basket analysis**. This problem deals with generation of
ftp.cise.ufl.edu/pub/faculty/ranka/Proceedings/p1.ps

LINK: Exploring Combinatorial Objects - Berry, Dean (1996) (Correct)
is known in the marketing community as **market basket analysis**. A system designed to explore the data
dimacs.rutgers.edu/pub/berryj/gd96.ps

Aspects of Network Visualization - Eick (1996) (Correct)
2 shows grocery purchasing patterns from a **market basket analysis**. The node sizes and colors show the
www.bell-labs.com/~eick/bibliography/1996/netvis_copyright.ps.gz

Try your query at: [Amazon](#) [Barnes & Noble](#) [Google \(RI\)](#) [Google \(Web\)](#) [CSB](#) [DBLP](#)

CiteSeer - citeseer.org - [Terms of Service](#) - [Privacy Policy](#) - Copyright © 1997-2002 [NEC Research Institute](#)

119 citations found. Retrieving documents...

Sergey Brin, Rajeev Motwani, and Craig Silverstein. *Beyond market baskets: Generalizing association rules to correlations*. In Proc. ACM SIGMOD, pages 265{ 276, 1997.

CiteSeer [Home/Search](#) [Document Details and Download](#) [Summary](#) [Related Articles](#) [Check](#)

This paper is cited in the following contexts:

[First 50 documents](#) [Next 50](#)

[Mining First-order Knowledge Bases for Association Rules - Jamil](#) (Correct)

....complexity of even the best known methods remains high. **While several ecient algorithms for association rule mining have been proposed [1, 6, 20, 15, 9, 19, 14, 21] overall ecieny is still a major issue, specially for other kinds of rule induction such as ratio rules [8] and chi square rules [4].** While many forms of rule inductions are interesting, association rules were found to be appealing because of their simplicity and intuitiveness. **In this paradigm, the rule mining process is divided into two distinct steps discovering frequent item sets and generating rules. There are**

....As future research, we plan to develop optimization techniques for mining queries that require non trivial look ahead and pruning techniques in aggregate functions. **The developments presented here also have other signi cant implications. For example, it is now possible to compute chi square rules [4] using the building blocks provided by our system.** Declarative computation of chi square rules, to our knowledge, has never been attempted for the many procedural concepts the computation of chi square method relies on. **In a separate work [2] we show that the counting method proposed in this paper**

Sergey Brin, Rajeev Motwani, and Craig Silverstein. *Beyond **market** baskets: Generalizing association rules to correlations*. In Proc. ACM SIGMOD, pages 265{ 276, 1997.

[Association Rule Mining on Remotely Sensed Imagery Using P-Trees - Ding \(2002\)](#) (1 citation) (Correct)

....first step, which is the generation of frequent itemsets [AS94] Having determined the frequent itemsets, the second step is very straightforward and provides few possibilities for improvement. **The reason is that **confidence** does not have any closure property while support has a downward property [BMS97].** By downward property, we mean that, if a set has a property, then all its subsets also have this property. **Support** is downward closed because of the fact that, if a set of items satisfies the minimum support, then all its subsets also satisfy the minimum support. **The downward closure property**

....study of interestingness measures for association rule patterns is given in [TK00] There are some critiques of the support **confidence** framework because this framework does not address some problems such as negative implications. **In addition, it may lead to misleading rules in some situations [BMS97, SBM98].** Table 2.3 gives a tea coffee example in the contingency table, where u represents the presence of an item and # its absence, and the numbers represent percentages of purchase. **Table 2.3. Contingency table of tea and coffee purchase**

	TEA	COFFEE	row sum
TEA	20	70	55
COFFEE	90	10	100
column sum	25	75	100

....

[Article contains additional citation context not shown here]

S. Brin, R. Motwani, and C. Silverstein, "*Beyond **Market** Baskets: Generalizing Association Rules to*

Correlations," Proceedings of the ACM SIGMOD, Tucson, AZ, May 1997, pp. 265-276.

Discovering Compact and Highly Discriminative Features or.. - Yu, Yang, Wang, Hah (Correct)

....(a) e) 50 100 No (a, b) d) 25 100 Yes (a, b, e) d) 25 100 Yes (f) d) 0 0 No (y) d) 100 100 Yes (B)
Examples of Association Rules Figure 1. **Transactions and Association Rules The revision of the Apriori algorithm adopting the chi squared test has been investigated [4].** This method suffers from generating too many uncorrelated rules because it still uses the support threshold [4]. S. Morishita suggested a scalable statistical pruning method by computing an upper bound of a statistical metric such as chi squared value, but the upper bound of the statistical

....100 Yes (B) Examples of Association Rules Figure 1. **Transactions and Association Rules The revision of the Apriori algorithm adopting the chi squared test has been investigated [4].** This method suffers from generating too many uncorrelated rules because it still uses the support threshold [4]. S. Morishita suggested a scalable statistical pruning method by computing an upper bound of a statistical metric such as chi squared value, but the upper bound of the statistical metric is only valid for binary feature set [15]. Correlation techniques have the following limitations: They

S. Brin, R. Motwani, and C. Silverstein. *Beyond market baskets: generalizing association rules to correlations.* In Proc. ACM SIGMOD International Conference on Management of Data, pages 265-276, Tucson, Arizona, 1997.

Using Association Rules for Product Assortment.. - Brijs, Swinnen.. (1999) (12 citations) (Correct)

.... cigarette paper [absolute sup = 291, conf = 0.82] **These rules demonstrate that whenever a customer buys cigarette paper, he/she also buys tobacco (confidence = 100) and that when a customer buys tobacco he will often also buy cigarette paper with it (confidence 82).** A more formal method [9] to assess the dependence between two or more products is of interest. Definition 5: Interest $s(X \rightarrow Y) = \frac{s(X \wedge Y)}{s(X)}$ The numerator $s(X \wedge Y)$ measures the observed frequency of the co-occurrence of the items in the antecedent (X) and the consequent (Y) of the rule. The denominator $s(X)$ is (Y)

Brin, S., Motwani, R., and Silverstein, C. *Beyond market baskets: generalizing association rules to correlations.* In Peckham, J., (ed.). Proceedings of the ACM SIGMOD Conference on Management of Data, 1997 (SIGMOD'97), 265-276.

Levelwise Search of Frequent Patterns with Counting.. - Bastide, Taouil.. (Correct)

....has been conducted on this topic. **The problem of mining frequent patterns arose first as a sub problem of mining association rules, but then it turned out that frequent patterns solve a variety of problems: mining sequential patterns [AS95] episodes [MTV97] association rules [AS94] correlations [BMS97, SBM98], multi dimensional patterns [KHC97, LSW97] maximal patterns [ZPOL97, LK98] and several other important knowledge discovery tasks [HPY00].** Since the complexity of this problem is exponential in the size of the binary database input relation and since this relation has to be scanned several times

S. Brin, R. Motwani, and C. Silverstein. *Beyond market baskets: Generalizing association rules to correlation.* Proc. SIGMOD conf., pp 265-276, May 1997.

Efficient Data Mining Based on Formal Concept Analysis - Stumme (Correct)

....many work has been conducted on this topic. The problem of mining frequent patterns arose first as a sub problem of mining association rules [1] but it then turned out to be present in a variety of problems [18] mining sequential patterns [3] episodes [26] association rules [2] correlations [10, 37], multi dimensional patterns [21, 22] maximal patterns [8, 53, 23] closed patterns [47, 31 33] Since the complexity of this problem is exponential in the size of the binary database input relation and since this relation has to be scanned several times during the process, efficient algorithms for

S. Brin, R. Motwani, and C. Silverstein. *Beyond **market** baskets: Generalizing association rules to correlation*. In Proc. ACM SIGMOD Int'l Conf. on Management of Data, pages 265-276, May 1997.

Supporting User Interaction for the Exploratory Mining of. - Mah (Correct)

....in association mining is developing parallel mining algorithms for finding association rules [12] 21] Other researchers are concerned with different issues; one recent debate is the appropriateness of using **confidence** to assess relationship or association. Brin, Motwani and Silverstein in [9] suggested that the dependence ratio or correlation between two sets are more appropriate to calculate relationships than **confidence**. The algorithm they proposed involved merging the frequent set generation and correlation calculation algorithms into one to increase the pruning power of the

....value in the output. 52 T is true T is false Sum of Row S is true sltl slt0 sl S is false s0tl s0t0 s0 Sum of Column tl tO n Table 3.3: Contingency Table for Association Rules 3.6. 2 **Correlation** Some researchers believe that correlation is a better measure of association than **confidence** [9]. Given this, the prototype allows the user to choose correlation as a relationship metric versus **confidence**. We believe that giving the user this choice makes the new exploratory model more flexible, redefining the definition of relationship in association mining to be any metric one sees fit

[Article contains additional citation context not shown here]

S. Brin, R. Motwani and C. Silverstein. *Beyond **market basket**: Generalizing association rules to correlations*. In Proc. ACM SIGMOD Conference, pages 265-276, 1997.

Parallel Formulations of Tree-Projection-Based Sequence.. - Guralnik, Karypis (Correct)

....those based on the **level** wise paradigm; nevertheless, they still require a substantial amount of time. A number of efficient and scalable parallel formulations have been developed for finding frequent itemsets and sequences that are based on the candidate generation and counting framework [3, 18, 22, 16, 4], both for shared and distributed memory parallel computers [2, 22, 17, 8, 25, 29, 20] However, the problem of parallelizing equivalence class based and projection based algorithms has received relatively little attention and existing parallel formulations for them have been targeted only toward

....[27] algorithm extended the Apriori like **level** wise mining method to find frequent patterns in sequential datasets. The basic **level** wise algorithm has been extended in a number of different ways leading to more efficient algorithms such as DHP [19, 18] Partition [22] SEAR and Spear [16] and DIC [4]. An entirely different approach for finding frequent itemsets and sequences are the equivalence class based algorithms Eclat [32] and SPADE [31] that break the large search space of frequent patterns into small and independent chunks and use a vertical database format that allows them to

S. Brin, R. Motwani, and C. Silverstein. *Beyond **market** baskets: Generalizing association rules to*

correlations. In Proc. of 1997.

Exploratory Mining via Constrained Frequent Set Queries - Ng, Lakshmanan, Hah, Mah (1999)
(10 citations) (Correct)

....exceed given thresholds. **While** such associations are useful, other notions of relationships may also be useful. **First**, there exist several significance metrics other than **confidence** that are equally meaningful. **For example, Brin et al. argue why correlation can be more useful in many circumstances [2].** Second, there may be separate criteria for selecting candidates for the antecedent and consequent of a rule. **For example, the user may want to find associations from sets of items to sets of types. Coming from different domains, the antecedent and consequent may call for different support**

S. Brin, R. Motwani, and C. Silverstein. *Beyond **market basket**: Generalizing association rules to correlations*. In Proc. 1997 ACM-SIGMOD, pp 265-276.

Finding Frequent Patterns Using Length-Decreasing Support.. - Seno, Karypis (Correct)

....algorithm extended the Apriori like **level** wise mining method to find frequent patterns in sequential databases. **The basic **level** wise algorithm has been extended in a number of different ways leading to more efficient algorithms such as DHP [14, 13] Partition [19] SEAR and Spear [12] and DIC [5].** An entirely different approach for finding frequent itemsets and sequences are the equivalence class based algorithms Eclat [26] and SPADE [24] that break the large search space of frequent patterns into small and independent chunks and use a vertical database format that allows them to

S. Brin, R. Motwani, and C. Silverstein. *Beyond **market** baskets: Generalizing association rules to correlations*. In Proc. of 1997 ACM-SIGMOD Int. Conf. on Management of Data, Tucson, Arizona, 1997.

STAMP: On Discovery of Statistically Important Pattern.. - Yang, Wang, Yu (Correct)

....was not fully taken into account by the multiple support model . In contrast, the generalized information gain metric would capture the difference of occurrences between B and C. 5.2. 3 **Statistically Significant Patterns There are much work in discovering statistically significant patterns [5, 18, 27].** All those work only takes into account the occurrence of a pattern in a sequence or a transaction. **However**, it does not assign any penalty if a pattern fails to be present when it is supposed to. **In** addition, all those work only discovers the significant patterns for the entire data set, and

S. Brin, R. Motwani, C. Silverstein. *Beyond **market** baskets: generalizing association rules to correlations*. Proc. ACM SIGMOD Conf. on Management of Data, 265-276, 1997.

Closed Set Based Discovery of Small Covers for Association.. - Pasquier, Bastide, Taouil (1999)
(9 citations) (Correct)

....according to the user preferences. **In** contrast, the second trend addresses the problem with an a priori vision, by attempting to minimize the number of exhibited rules. **In** [14, 28] information about taxonomies are used to define criteria of interest which apply for pruning redundant rules. **In** [7, 25], **statistical measures such as Pearson s correlation or the chi squared test are used instead of the **confidence** measure.** 1.2 Contribution: an Overview The approach presented in this paper belongs to the second trend since it aims to extract not all possible rules but a sub set called small cover

....specified patterns. **Information** in taxonomies associated with the dataset can also be integrated in the process as proposed in [14, 28] for extracting bases for generalized (multi **level**) association rules. **Integrating item constraints and statistical measures, such as described in [5, 22, 29] and [7, 25] respectively, in the generation of bases requires further work.** Functional and approximate dependencies Algorithms presented in this paper can be adapted to generate bases for functional and approximate dependencies. In [15, 20] such bases and algorithms for generating them were proposed.

S. Brin, R. Motwani, and C. Silverstein. *Beyond **market** baskets: Generalizing association rules to correlation*. Proc. of the ACM SIGMOD Conference, pages 265-276, May 1997

Parallel Formulations of Tree-Projection-Based Sequence.. - Guralnik, Karypis (Correct)

....those based on the **level** wise paradigm; nevertheless, they still require a substantial amount of time. A **number of efficient and scalable parallel formulations have been developed for finding frequent itemsets and sequences that are based on the candidate generation and counting framework [3, 18, 22, 16, 4], both for shared and distributed memory parallel computers [2, 22, 17, 8, 25, 29, 20]** However, the problem of parallelizing equivalence class based and projection based algorithms has received relatively little attention and existing parallel formulations for them have been targeted only toward

....[27] algorithm extended the Apriori like **level** wise mining method to find frequent patterns in sequential datasets. **The basic **level** wise algorithm has been extended in a number of different ways leading to more efficient algorithms such as DHP [19, 18] Partition [22] SEAR and Spear [16] and DIC [4].** An entirely different approach for finding frequent itemsets and sequences are the equivalence class based algorithms Eclat [32] and SPADE [31] that break the large search space of frequent patterns into small and independent chunks and use a vertical database format that allows them to

S. Brin, R. Motwani, and C. Silverstein. *Beyond **market** baskets: Generalizing association rules to correlations*. In Proc. of 1997 ACM-SIGMOD Int. Conf. on Management of Data, Tucson, Arizona, 1997.

Optimization of Constrained Frequent Set Queries with.. - Lakshmanan, Ng, Hah (1998) (23 citations) (Correct)

.... **group includes studies that go beyond the initial notion of association rules to other kinds of mined rules, e.g. multi **level** rules [8, 21] quantitative and multi dimensional rules [22, 7, 14, 10] rules with item constraints [23] mining long patterns [3] correlations and causal structures [4, 20], ratio rules [12] etc.** Recently it has been recognized that the integration of data mining technologies with database management systems is of crucial importance [5] Furthermore, it has been argued that the fundamental distinction of a data mining system from a statistical **analysis** program or

S. Brin, R. Motwani, and C. Silverstein. *Beyond **market** basket: Generalizing association rules to correlations*. In Proc. 1997 ACM-SIGMOD, pp 265-276.

Finding Frequent Patterns Using Length-Decreasing Support.. - Seno, Karypis (Correct)

....algorithm extended the Apriori like **level** wise mining method to find frequent patterns in sequential databases. **The basic **level** wise algorithm has been extended in a number of different ways leading to more efficient algorithms such as DHP [14, 13] Partition [19] SEAR and Spear [12] and DIC [5].** An entirely different approach for finding frequent itemsets and sequences are the equivalence class

based algorithms Eclat [26] and SPADE [24] that break the large search space of frequent patterns into small and independent chunks and use a vertical database format that allows them to

S. Brin, R. Motwani, and C. Silverstein. *Beyond **market** baskets: Generalizing association rules to correlations*. In Proc. of 1997.

Exploratory Mining and Pruning Optimizations of.. - Ng, Lakshmanan, Pang.. (1998) (85 citations)
(Correct)

....exceed given thresholds. While such associations are useful, other notions of relationships may also be useful. First, there exist several significance metrics other than **confidence** that are equally meaningful. For example, Brin et al. argue why correlation can be more useful in many circumstances [5]. Second, there may be separate criteria for selecting candidates for the antecedent and consequent of a rule. For example, the user may want to find associations from sets of items to sets of types. The rule peps snacks is an instance of such an association, meaning that customers often buy the

....Phase II, the user can specify the desired significance metric, and can give different conditions that must be satisfied by the antecedent and consequent of the relationships to be formed. There are already several proposals in the literature that make the notion of associations less rigid [5, 7, 9, 12, 14, 21]. We are not proposing another here. Instead, we are proposing an architecture that allows many of those alternative notions to co exist, and that permits the user to choose whatever is appropriate for the application. 2 Architecture Figure 1 shows a two phase architecture for exploratory

S. Brin, R. Motwani, and C. Silverstein. *Beyond mar- ket **basket**: Generalizing association rules to correlations*. SIGMOD 97, pp 265-276.

Frequency-Based Views to Pattern Collections - Mielikäinen (2003) (Correct)

....X Y where X and Y are subsets of R. The most popular interestingness measure for an association rule X Y is its accuracy (or **confidence**) which is defined as $\text{acc}(X Y, d) = \frac{\text{fr}(X Y, d)}{\text{fr}(X, d)}$. Also several other classes of patterns and measures of interestingness have been studied (see e.g. [4, 6, 9, 13, 14, 15, 27, 32, 33, 36, 37, 43, 44, 47, 48]) It is not always easy to define an interestingness measure # in such a way that there would be a threshold value # such that #(p) # for almost all interesting patterns p and for only very few uninteresting ones. One way to augment the interestingness measure is to define additional

S. Brin, R. Motwani, and C. Silverstein, *Beyond **market** baskets: Generalizing association rules to correlations*, in SIGMOD 1997.

TAR: Temporal Association Rules on Evolving Numerical Attributes - Wang, Yang, Muntz (2001)
(Correct)

...., is $t \in \text{Gammam } 1 \leq j \leq N (P_i; W(j; m))$ where $N (P_i; W(j; m))$ is the number of object histories which follow P_i on window $W(j; m)$ 3.1.2 Strength Different methods can be used to capture the degree of nonindependence. In this paper, we use a metric that is similar to interest defined in [4] to measure the strength of a temporal association rule. Definition 3.3 Given a temporal association rule $R : X (Y$ and a sequence of : $S_1 ; S_2 ; S_t$, the strength of the rule is $\text{Support}(X Y; \Omega \text{Gamma}) \text{Support}(X; \Omega \text{Gamma}) \text{ThetaSupport}(Y; \Omega \text{Gamma})$. 3.1.3 Density Since

S. Brin, R. Motwani, C. Silverstein. *Beyond **market** baskets: generalizing association rules to*

correlations. Proc. ACM SIGMOD Conf. on Management of Data, 265-276, 1997.

InfoMiner: Mining Surprising Periodic Patterns - Jiong Yang Jiyang (Correct)

....(i.e. 1) towards its significance, regardless of its likelihood of occurrence. **Intuitively, the assessment of significance of a pattern in a sequence should take into account the expectation of pattern occurrence (according to some prior knowledge)** Recently, many research has been proposed [1, 3, 5, 6, 8, 9, 10, 11, 12, 15] towards this objective. We will furnish an overview in the next section. In this paper, a new model is proposed to characterize the class of so called surprising patterns (instead of frequent patterns) We will see that our model not only has a solid theoretical foundation but also allows an

....not fully taken into account by the multiple support model. **In contrast, the information gain metric proposed in this paper would capture the difference of occurrences between B and C. Mining patterns that are statistically significant (rather than frequent) becomes a popular topic. Brin et al. [3] first introduced the concept of correlation and it was shown that in many applications the correlation measurement can reveal some very important patterns.** The Chi squared test was used to test the correlation among items. **Instead** of explicitly enumerating all correlated itemsets, the border

[Article contains additional citation context not shown here]

S. Brin, R. Motwani, C. Silverstein. *Beyond **market** baskets: generalizing association rules to correlations*. Proc. ACM SIGMOD Conf. on Management of Data, 265-276, 1997.

Generating Dual-Bounded Hypergraphs - Boros, Elbassioni, Gurvich.. (2002) (Correct)

....respectively, to the minimal infrequent and maximal frequent sets for D. **The generation of (maximal) frequent sets of a given binary **matrix** is an important task of knowledge discovery and data mining, e.g. it is used for mining association rules [7, 31, 52, 53, 56, 57, 70] correlations [20], sequential patterns [2] episodes [54] emerging patterns [25] and appears in many other applications.** Most practical procedures to generate frequent sets are based on the anti monotone Apriori heuristic (see [1] and build frequent sets in a bottom up way, running in time proportional to the

S. Brin, R. Motwani, and C. Silverstein, *Beyond **market basket**: Generalizing association rules to correlations*, Proc. the 1997 ACM-SIGMOD Conference on Management of Data, pp. 265-276. -- 21 --

OSSM: A Segmentation Approach to Optimize **Frequency** Counting - Leung (Correct)

....find cardinalities of subgroups or significance of deviations, etc. **Typically, the patterns, whose frequencies are needed, are conjunctions of atomic patterns. A prime example is given by the frequent set concept underlying association rules [2, 3] Moreover, the patterns defined for correlation [6, 7], causality [18] sequential patterns [4] episodes [13] constrained frequent sets [11, 14, 19] long patterns [1, 5] closed sets [16] and many other important data mining tasks have the same basic form.** In all these cases, we have instances of the following abstract problem. **Given** a collection

S. Brin, R. Motwani, and C. Silverstein. *Beyond **market basket**: Generalizing association rules to correlations*. In Proc. SIGMOD 1997, pp 265--276.

Optimization of Association Rule Mining Queries - Jeudy, Boulicaut (2002) (Correct)

....the frequent itemsets and their frequencies are needed. **This is an important point, and we consider in the rest of this paper that the generation of association rules does not need to access the transactional database (it is still the case when using other objective measures such as the conviction [9] or the J measure [25] However, we allow any other constraint on the association rules and we do not require the occurrence of the frequency and or confidence constraints.** Given an association rule constraint C, let us study different strategies to support constrained association rule mining

S. Brin, R. Motwani, and C. Silverstein. *Beyond **market** baskets: Generalizing association rules to correlations*. In J. M. Peckman, editor, Proceedings of ACM SIGMOD Conference on Management of Data (SIGMOD '97), pages 265-276, Tucson, AZ, May 1997. ACM.

Data Mining of Association Rules and the Process of.. - Hipp, Güntzer.. (Correct)

....a counter is set up and the algorithm then passes over the complete database of transactions. **Whenever** a transactions contains one of the candidates its counter is incremented. **Efficiently** looking up candidates in transactions requires specialized data structures, e.g. hashtrees or prefix trees, c.f. [3, 6]. Alternatively the support values of candidates can be determined indirectly by set intersections. **For** that purpose so called transaction sets are employed. **The** transaction set $X:tids$ of an itemset X is defined as the set of all transactions this itemset is contained in: $X:tids = \{t \in D \mid X \subseteq t\}$

....Partition [26] combines the breadth first search of Apriori with determining the support values of the candidates indirectly by set intersections. **In** order to be able to keep all necessary transaction sets comfortably in main memory the database typically needs to be partitioned. **The algorithm DIC [6] enhances Apriori by relaxing the strict separation between candidate generation and counting the candidates.** Already during passing over the transactions new candidates are generated and added to the set of candidates on the fly. **This** helps to significantly reduce the total number of the

S. Brin, R. Motwani, and C. Silverstein. *Beyond **market** baskets: Generalizing association rules to correlations*. In Proceedings of the ACM SIGMOD International Conference on Management of Data (ACM SIGMOD '97), pages 265-276, 1997.

Generating Dual-Bounded Hypergraphs - Boros, Elbassioni, Gurvich.. (2002) (Correct)

....respectively, to the minimal infrequent and maximal frequent sets for D . **The generation of (maximal) frequent sets of a given binary **matrix** is an important task of knowledge discovery and data mining, e.g. it is used for mining association rules [7, 32, 53, 54, 58, 59, 74] correlations [20], sequential patterns [2] episodes [55] emerging patterns [25] and appears in many other applications.** Most practical procedures to generate frequent sets are based on the anti monotone Apriori heuristic (see [1] and build frequent sets in a bottom up way, running in time proportional to

S. Brin, R. Motwani, and C. Silverstein(1997). *Beyond **market basket**: Generalizing association rules to correlations*. Proc. ACM-SIGMOD Conference on Management of Data, 265--276.

Profiling High **Frequency** Accident Locations Using.. - Geurts, Wets, Brijs.. (Correct)

No context found.

Brin, S., Motwani, R. and C. Silverstein. *Beyond **market** baskets: generalizing association rules to correlations*. In Proceedings of the ACM SIGMOD Conference on Management of Data, Tucson,